Knowledge Servers – Applications of Artificial Intelligence to Advanced Space Information Systems

PETER FRIEDLAND

AI RESEARCH BRANCH, MAIL STOP 244-17
NASA AMES RESEARCH CENTER
MOFFETT FIELD, CA 94035

NASA Ames Research Center
Artificial Intelligence Research Branch

Technical Report RIA-88-02-01-01

February, 1988

Knowledge Servers -- Applications of Artificial Intelligence to Advanced Space Information Systems

Peter Friedland
Artificial Intelligence Research Branch
NASA Ames Research Center
Moffett Field, California 94035

Abstract

We have begun a transition from passive information systems which act only to facilitate the storage and retrieval of stereotyped data to far more active and responsive systems which can deal with widely differing forms of human knowledge. Edward Feigenbaum has coined the term "knowledge servers" to describe this next generation of active information management systems. Among the functions of a knowledge server will be: the ability to store enormous varieties of knowledge; the ability to determine, through natural discourse, the needs of its users; the ability to summarize and pursue complex relationships in its knowledge; the ability to test and critique user hypotheses and suggest previously unseen connections resulting from those hypotheses; and the ability to communicate and collaborate with other autonomous knowledge servers. Because of the complexity and variety of information relevant to future major space missions like space station, these missions will act as a driving force and testbed for the knowledge server concept.

1. Introduction

In the early days of artificial intelligence research it was thought that the key to powerful problem-solving behavior was in the reasoning process, that machines needed to utilize complex, general-purpose "inference engines" in order to emulate human performance in real world domains. This approach led to much frustration and many blind alleys until the discovery in the early 1970's (through the development of such early "expert systems" as DENDRAL, MACSYMA, and MYCIN) that it was in the knowledge, not the reasoning processes, where the power lay. This discovery was tested through the creation of thousands of such knowledge-based systems in the 1970's and 1980's. Many were capable of near-expert performance in complex domains using very simple inference mechanisms; all acted as confirmation of the knowledge is power thesis. This thesis has now become formalized as The Knowledge Principle: that a computational system can exhibit a high level of competence in understanding and action primarily because of the domain specific knowledge it contains.

An importance consequence of The Knowledge Principle is that intelligent systems must contain large amounts of knowledge (in its many and varied forms) about the domains in which they perform. Therefore, the pragmatics of building such a system dictate the need for the ability to acquire knowledge, to represent knowledge, and to manipulate knowledge; these topics have formed the core of artificial intelligence research since 1970.

Space Information Systems in the Space Station Era and beyond will be among the largest knowledge-based systems. They will need to store and effectively retrieve both structural and functional knowledge of artifacts as large as the Space Station, and will need to relate that knowledge to queries from a wide variety of other intelligent agents, both human and machine. The remainder of this paper will discuss current bottlenecks in building such comprehensive systems and will analyze the major research topics that will lead to the solution of those bottleneck problems. The paper will conclude with a vision of the "dream" space information system--the knowledge server for Space Station.

2. Current Bottlenecks

Knowledge-based systems have already been built that exhibit expert-level performance on significant problems. However, the amount of knowledge available to these systems is at least two orders of magnitude less than a comprehensive knowledge base relating to Space Station would require. Current representation mechanisms employ at least one of three basic structures: statements in the predicate calculus (logic), if-then rules, and collections of attribute-value pairs known as frames. The best current systems utilize all three structures in a hybrid approach. A typical "large" present-day knowledge-based system might contain up to a few thousand rules and/or statements in the predicate calculus, along with a few hundred frames.

There are four major bottlenecks in scaling systems up to the size we will require in the future:

- Knowledge Acquisition--Getting the knowledge into the system.

- Knowledge Consistency and Completeness--Combining knowledge from many, potentially conflicting sources in many different forms and determining if enough knowledge has been stored for functional adequacy.

- Large Knowledge Base Manipulation--Determining if and how current methodologies for knowledge manipulation scale up to the demands of far larger systems.

- Interface Technology--Getting the knowledge out in a form that humans or other intelligent agents can understand and utilize effectively.

2.1. Knowledge Acquisition

Developing the knowledge base for a current-generation system currently requires the participation of a human "knowledge engineer" who acts as the bridge between domain expertise and computational representations. While knowledge engineers have become highly skilled at their task, they cannot avoid the introduction of significant delay and loss of accuracy into the knowledge base construction project. Major research efforts are underway to eliminate this intermediary by means of both automatic knowledge acquisition (by both inductive and deductive means) and by semi-automatic knowledge acquisition where domain experts act as their own knowledge engineers.

In a typical knowledge-based system project, knowledge acquisition can take anywhere from 50% to 90% of the time and effort; for systems two orders of magnitude larger than current systems, clearly, better mechanisms are necessary. In addition, since systems such as those which are to function as information systems for Space Station must have dynamic abilities to grow and improve their behavior, knowledge acquisition cannot be viewed as a one-time task, but a continuing one throughout the life of the system.

2.2. Knowledge Consistency and Completeness

The construction of a current knowledge systms typically involves interaction between a single (or at most a few) expert and a small team of knowledge engineers. But for the large-scale systems of the future, no single (or even small group) expert can possibly know more than a small fraction of the required knowledge. In addition, much of the expertise will come from other than human sources; it may arise directly from CAD systems during design or it may come during automated construction and testing processes. A major problem occurs in combining expertise from multiple sources. There may be disagreements both over facts (tolerances, expected lifetimes, etc.) and over heuristics (why certain design decisions were made, where components are likely to fail first, etc.). But even where disagreements are not critical, the "language" used to describe knowledge can be expected to vary significantly from source to source. Mechanisms are needed to ensure that large-scale knowledge systems are not highly dependent upon either the source or order of information entry.

A second issue arises of knowledge completeness for any significantly sized system, i.e. when is there enough knowledge to do an adequate job of problem-solving. In a sense, this problem is even more difficult than the knowledge consistency bottleneck since any notion of completeness involves "deep" knowledge of both the domain and the desired functional attributes of the knowledge system. Current systems have none of this "self-awareness," and the only present-day answers are for completeness in the formal sense for logic-based knowledge systems. The problem of knowledge completeness should be thought of as dynamic, with criteria changing during the evolutionary growth of missions like Space Station.

2.3. Large Knowledge Base Manipulation

The artificial intelligence research community has developed, over the last decade, adequate methods for reasoning with our present knowledge bases. These inference methods tend to be simple (following the Knowledge Principle) mechanisms like forward and backward chaining of rules, theorem proving in the predicate calculus, and object-oriented programming (attached procedures) for frame representations. However, we do not yet know whether our current methods will scale-up to the several orders-of-magnitude larger systems of the near future. Perhaps hardware solutions will come to the rescue and enable us to continue to rely on our present inferential methods, but it is equally likely that future problem knowledge base size and complexity will spur research into a second generation of more comprehensive and efficient reasoning methods. At any rate, much empirical experimentation will be required.

2.4. Interface Technology

Current systems are able to communicate actions and (perhaps as important) reasons for actions in a variety of formats ranging from graphical to textual. But even the best of present knowledge and information systems do not employ more than a tiny subset of true natural language nor do they employ, in any real sense, that most common of human interaction mechanisms, speech. As knowledge systems become larger, more comprehensive, and more ubiquitous, mechanisms for smooth interaction with other intelligent agents (particularly, but not exclusively humans) will become even more important.

3. Brittleness of Current Systems

A characteristic of all current computational systems (whether traditional or knowledge-based) is extreme brittleness at the edge of their performance. Even the most "user-friendly" computer programs are trivially easy to confuse with the wrong syntactic or semantic input. In the case of knowledge-based systems, even those that perform at extremely high levels of expertise within narrowly limited domains degrade rapidly to incompetence at the limits. Part of the problem comes from a total lack of self-awareness on the part of computational systems. Humans (often) know when they are near or beyond the bounds of their expertise and switch to more generic knowledge and problem-solving mechanisms or call in other, more expert consultants for advice. Part of the problem comes simply from the enormous gap between the knowledge of the least competent human and the most competent knowledge-based system. And finally, part of the problem arises from the fact that most operational knowledge-based systems are primarily heuristic or experiential; they rely solely upon compiled, mostly anecdotal experience, without a detailed model of structure and function to fall back on.

For space information systems, the brittleness problem is critical. We cannot have an information resource that simply gives up when precise, encapsulated expertise in unavailable. But, on the other hand, a system which pretends to expertise when its information is only a very rough guess is equally dangerous. The next three sections of this paper discuss potential solutions to the brittleness problem for knowledge systems.

4. Common Sense and Analogical Reasoning

In situations where our particular expertise does not apply, we humans rely upon two, probably related, forms of reasoning. The first we label "common sense" and seems to encompass the generic knowledge which we can apply to a broad range of situations and actions. Both philosophers and computer scientists have pondered the exact nature of common sense knowledge; current thinking within the artificial intelligence community is that common sense comes from knowing a little about a lot and recognizing, from vast and continual experience with our environment, relevance and intereonnectivity of knowledge. Dr. Douglas Lenat, of the Microelectronics and Computer Consortium (MCC), is currently testing this thesis in his ten-year CYC Project by attempting to build a knowledge base which contains everything in a desk encyclopedia and then determining if the resultant knowledge system responds gracefully over an extremely broad range of queries and problems.

The second mechanism is called analogical reasoning. Analogies are simply a method for relating common properties of seemingly different situations. Sometimes they are very similar situations; we may know what to expect in a Cambodian restaurant because we have already been to a Thai one. Sometimes, however, they allow us to reason about wildly different situations; understanding genetic regulation by thinking of what can slow and divert a locomotive on railroad tracks is an example of that. For all analogies, the reasoning process consists of two steps: pick a potential analogous situation, and extract out the common items. While the process sounds simple, little real progress has been made in building computational systems which reason by analogy. Many believe the lack of progress is caused by the fact that to effectively utilize analogies, as discussed for common sense reasoning above, the system needs to know a little about a lot; indeed the ability to form analogical reasoning is a major test of the CYC system. In any event, the ability to effectively analogize will add a major dimension of reasoning power to any future information management system.

5. Learning

Perhaps more than any other single characteristic, the ability to learn is what separates human problem-solvers from current computational ones. We continually improve and expand our own internal "knowledge bases" by a variety of mechanisms including explicitly from teachers, colleagues, and written material, and implicitly by experimentation and discovery from our environment. We would consider a fellow human impossibly stupid if the identical mistake was endlessly repeated, but almost all of our computational partners fall into this category. As our knowledge systems grow larger and larger, the addition of learning mechanisms will be the only hope for continuing to manage complexity.

Space Station is designed to continue and evolve over a period of at least 35 years. Even if we imagine a perfectly knowledgable information system at the initial configuration of Space Station, such a system will not be useful for very long as systems change by design or accident and missions change as we better understand the unique utility of the facility. But if each change requires specific, conscious, human-mediated changes to the information system, then humans will need to continue to be global experts in every piece of knowledge to recertify consistency of the knowledge base.

Learning mechanisms bring their own complexity to information mangement. As important as the ability to learn is the knowledge of when learning is worthwhile; clearly, every minor "discovery" need not be stored away for later use or even the large knowledge systems of the future will become overwhelmed. In addition, in the validation of computational systems, a decision must be made as to whether a learning mechanism can be validated as safe and reliable or whether each change, no matter how minor, to the knowledge system must be independently revalidated.

6. Synergism Among Intelligent Agents

Humans do not operate in an information vacuum; even experts cooperate with other humans of varying degrees of sophistication. Much of our power as problem-solvers comes from knowing when to seek out help in areas in which we are not expert. A powerful characteristic of future knowledge systems must be the ability to seek out answers from other sources, both human and machine, on the information network. Humans find that the synergism of multiple expertise causes the whole to often be much more than the sum of its parts for problem-solving. We have begun to develop the hardware mechanisms (the ethernet for distributed systems, for example) to allow such cooperation. However, the technology to know when and where to seek synergism is still well beyond the state-of-the-art. Researchers in artificial intelligence have begun to explore mechanisms, spanning a continuum from hierarchical to distributed, for achieving such intelligent behavior among intelligent computational agents, but the work is still in its infancy. In addition, much research is needed in such areas as communication of intent and abilities among intelligent agents.

7. The Dream for Space Information Systems

This paper has attempted to sketch the current state of knowledge systems along with the areas of research necessary to turn them into the kind of information utilities that will be truly useful on missions like Space Station. Let us imagine the properties of this new system, called a Knowledge Server, on our information network for Space Station.

- The Knowledge Server can determine, through natural discourse, the needs and desires (both overt and hidden) of its users. It maintains a memory of abilities and preferences of those it has served in the past, and can determine a reasonable model of the intents and abilities of new users. It is able to present answers in a form that is suitable to each particular user, whether astronaut or ground controller, specialist or generalist.

- The Knowledge Server can explain how it arrived at any particular answer and can justify its reasoning in response to user questions.

- The Knowledge Server can summarize and pursue complex relationships in its knowledge either when requested or when it believes such analysis will add to its understanding. It will actively seek out ways of improving its overall knowledge of Space Station and point out unseen and potential flaws in subsystems.

- The Knowledge Server can test and critique user hypotheses, working either autonomously or in close cooperation with a human user.

- The Knowledge Server can communicate and collaborate with other autonomous knowledge servers, both human and machine.

- The Knowledge Server can serve as an archival respository of all decisions, actions, and events that affect Space Station, and ensure that all relevant parties are informed as necessary about such knowledge.

Clearly, achieving the above abilities will be a non-trivial task; indeed, many view the research necessary as forming the core of work in artificial intelligence over the next several decades. However, building such Knowledge Servers to support the information needs of future space missions may be essential as our tasks grow increasingly complex and long-lived.

Acknowledgments

This article was inspired by Professor Edward Feigenbaum's keynote speech to the MEDINFO-86 conference, entitled "Knowledge Processing: From File Servers to Knowledge Servers." While the context of the speech was medical information processing, many of the ideas are transferable to space information processing. The author thanks Professor Feigenbaum for permission to use his speech as a starting point for his own thinking and writing on the matter.

Table of Contents

RIA-87-11-16-6

Automatic Bayesian Induction of Classes

PETER CHEESEMAN, JAMES KELLY, MATTHEW SELF, AND JOHN STUTZ — November 1987

This paper describes a criterion, based on Bayes' theorem, that defines the optimal set of classes (a classification) for a given set of examples. This criterion does not require that the number of classes be specified in advance; this is determined by the data. Tutored learning and probabilistic prediction in particular cases are an important indirect result of optimal class discovery. Extensions to the basic class induction program include the ability to combine category and real-valued data, hierarchical classes, independent classifications and deciding for each class which attributes are relevant.

RIA-88-02-01-01

Knowledge Servers - Applications of Artificial Intelligence to Advanced Space Information Systems

PETER FRIEDLAND — February 1988

We have begun a transition from passive information systems which act only to facilitate the storage and retrieval of stereotyped data to far more active and responsive systems which can deal with widely differing forms of human knowledge. Edward Feigenbaum has coined the term 'knowledge servers' to describe this next generation of active information management systems. Among the functions of a knowledge server will be: the ability to store enormous varieties of knowledge; the ability to determine, through natural discourse, the needs of its users; the ability to summarize and pursue complex relationships in its knowledge; the ability to test and critique user hypotheses and suggest previously unseen connections resulting from those hypotheses; and the ability to communicate and collaborate with other autonomous knowledge servers. Because of complexity and variety of information relevant to future major space missions like space station, these missions will act as a driving force and testbed for the knowledge server concept.

RIA-88-04-01-4

AutoClass: A Bayesian Classification System

PETER CHEESEMAN, JAMES KELLY, MATTHEW SELF, JOHN STUTZ, WILLIAM TAYLOR, AND DON FREEMAN — April 1988

This paper describes AutoClass II, a program for automatically discovering (inducing) classes from a database, based on a Bayesian statistical technique which automatically determines the most probable number of classes, their probabilistic descriptions, and the probability that each object is a member of each class. AutoClass has been tested on several large, real databases and has discovered previously unsuspected classes. There is no doubt that these classes represent new phenomena.

RIA-88-04-12-0

Learning by Making Models

PHILIP LAIRD — April 1988

We propose a theory of learning from unclassified data. The learning problem is that of finding the parameters of a stochastic process that best describes the incoming data stream. Special attention is given to the efficiency of the learning process, similar to Valiant's theory of supervised learning, and in contrast to conventional pattern recognition approaches. Illustrative domains are constructed and analyzed.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
Dates attached

3. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE

Titles/Authors - Attached

5. FUNDING NUMBERS

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Code FIA - Artificial Intelligence Research Branch
Information Sciences Division

8. PERFORMING ORGANIZATION REPORT NUMBER

Attached

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Nasa/Ames Research Center
Moffett Field, CA. 94035-1000

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Available for Public Distribution

Peter Friedland 5/14/92 BRANCH CHIEF

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Abstracts ATTACHED

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT

18. SECURITY CLASSIFICATION OF THIS PAGE

19. SECURITY CLASSIFICATION OF ABSTRACT

20. LIMITATION OF ABSTRACT